arXiv:1606.00546v1 [stat.AP] 2 Jun 2016

Forecasting Wind Power – Modeling Periodic and Non-linear Effects Under Conditional Heteroscedasticity

Florian Ziel

European University Viadrina, Chair for Finance and Capital Market Theory, Große Scharrnstraße 59, 15230 Frankfurt (Oder), Germany, Tel. +49 (0)335 5534 2986, E-Mail: ziel@europa-uni.de.

Carsten Croonenbroeck

University of Rostock, Faculty for Environmental Science, Justus-von-Liebig-Weg 2, 18059 Rostock, Germany, Tel. +49 (0)381 498 3267, E-Mail: carsten.croonenbroeck@uni-rostock.de.

Daniel Ambach

European University Viadrina, Chair for Quantitative Methods, esp. Statistics, Große Scharrnstraße 59, 15230 Frankfurt (Oder), Germany, Tel. + 49 (0)335 5534 2983, E-Mail: ambach@europa-uni.de.

Abstract

In this article we present an approach that enables joint wind speed and wind power forecasts for a wind park. We combine a multivariate seasonal time varying threshold autoregressive moving average (TVARMA) model with a power threshold generalized autoregressive conditional heteroscedastic (power-TGARCH) model. The modeling framework incorporates diurnal and annual periodicity modeling by periodic B-splines, conditional heteroscedasticity and a complex autoregressive structure with non-linear impacts. In contrast to usually time-consuming estimation approaches as likelihood estimation, we apply a high-dimensional shrinkage technique. We utilize an iteratively re-weighted least absolute shrinkage and selection operator (lasso) technique. It allows for conditional heteroscedasticity, provides fast computing times and guarantees a parsimonious and regularized specification, even though the parameter space may be vast. We are able to show that our approach provides accurate forecasts of wind power at a turbine-specific level for forecasting horizons of up to 48 hours (short- to medium-term forecasts).

Keywords: Renewable Energy, Wind Speed, Wind Power, Heteroscedasticity, Stochastic Modeling, Lasso
JEL: C13, C32, C53, Q47

1. Introduction

Wind power is on the verge of becoming the most important source of electricity in many countries worldwide. Berkhout et al. (2013) argue that wind power is the most emergent renewable power source with a growth rate of 30% per year. However, the technology still has a few challenges to master. In contrast to conventional power, wind power production is non-deterministic and highly volatile. To make efficient contracts at the energy pools, accurate forecasts of wind power production have to be available. Lei et al. (2009) as well as Soman et al. (2010) provide a time-scale classification of wind power and wind speed prediction models. Longer-term forecasts at horizons of two days up to one week for, e.g., decisions with respect to the required energy reserves and the maintenance scheduling, are based on meteorological or recently developed hybrid structure models.

Match-making at the energy markets, i.e. trading at the usual day-ahead markets common to most energy pools, requires predictions at forecasting horizons of up to 48 hours, at most, dependent on the designated contract market. Forecasts for this medium- to long-term scenario are usually based on stochastic modeling, on artificial intelligence models or specific neural networks. For instance, Cadenas and Rivera (2009), Cao et al. (2012) and Azad et al. (2014) use them. Amjady et al. (2011) use ridgelet neural networks which possess ridge functions as activators for their hidden nodes to provide forecasts of the aggregated wind power output of a wind farm. Bhaskar and Singh (2012) take a statistical approach which does not use numerical weather predictions. They use a wavelet decomposition of their wind speed time series and an adaptive wavelet neural network. After transformation, they transfer the wind speed predictions by using a feed-forward neural network into wind power forecasts. Liu et al. (2014) propose a hybrid model which combines inputs selected by deep quantitative analysis, wavelet transform, genetic algorithm and support vector machines. Another wavelet support vector machine approach is used by Zeng and Qiao (2012) to perform wind power predictions. Zhou et al. (2013) apply a probabilistic kernel density forecasting model with a quantile-copula estimator to perform wind power forecasts. They evaluate the model by using a power system in Illinois and compare several scheduling strategies. Haque et al. (2014) provide a new hybrid intelligent algorithm for wind power predictions that uses a combination of wavelet transform and fuzzy network methods.

For match-making, stochastic forecasting approaches like the one presented

in this paper benefit from modeling the persistence of wind power, its periodic structure and its direct dependence on wind speed. Thus, wind speed itself is usually predicted in an entirely stochastic setting, while also, numerical weather predictions (NWPs) can be employed, if available. Powerful statistical models return reliable forecasts of wind power for short- to medium-term scenarios and are widely established, like the Wind Power Prediction Tool (WPPT) by Nielsen et al. (2007), its recent generalization, GWPPT, by Croonenbroeck and Dahl (2014), or the spatial GWPPT by Croonenbroeck and Ambach (2015). However, the class of statistical approaches also incorporates autoregressive (AR), autoregressive moving average (ARMA) and autoregressive fractionally integrated moving average (ARFIMA) models. Kavasseri and Seetharaman (2009) discuss these models in details.

The literature on contributed models for wind power and wind speed forecasts is vast, Jung and Broadwater (2014) as well as Tascikaraoglu and Uzunoglu (2014) provide an up-to-date overview. Most models have several drawbacks: One problem is that stochastic wind power prediction models require wind speed forecasts in the first step. Several models for wind speed forecasting are available, as provided by Zhu et al. (2014), Ambach and Schmid (2015) or Shukur and Lee (2015). In the second step, these predictions are transformed into forecasts of wind power, as shown by, e.g., Azad et al. (2014). Most of the models do not provide conjoint wind power and wind speed predictions. Many models, e.g. the aforementioned WPPT class models, utilize wind speed as a quadratic regressor for wind power, although the theoretical non-linear relationship is usually described by a cubic function. The reason for this is to be found in the physical limitation of the turbine, i.e. the upper bound of producible wind power.

The long memory structure of usual turbine specific wind speed and wind power data suggests a diurnal and an annual periodic behavior. Several contributions illustrate this periodic or cyclic behavior, as, e.g., Carapellucci and Giordano (2013), Silva et al. (2016), Scholz et al. (2014) and Ambach and Schmid (2015). Ambach (2015) focuses on annual periodic effects.

Periodic effects may change over time, which is usually not considered. Ambach (2015) and Ambach and Croonenbroeck (2015) incorporate seasonal interactions in wind speed time series by annual and diurnal basis functions. Thereby, they capture the annual change of a daily period. This effect is basically driven by the fact that the length of the nights changes over the year: On the northern hemisphere, there are longer nights during the winter than during the summer. Indeed, it is observable that the diurnal periodicity is

varying over the year. Moreover, evidence suggests that such periodicities are also observable within the wind power data (see, e.g., Nielsen et al., 2007). Consequently, our new forecasting model includes all common stochastic modeling features, but in addition, it overcomes the aforementioned drawbacks. The main advantage of our approach is related to the fact that we are able to produce wind speed and power forecasts at the same time with one model. The periodic behavior of the day and the year is modeled by periodic B-splines. Furthermore, we capture interaction between both seasonalities, as the diurnal impact may change over the year. Thus, we allow for periodic changes in the parameters to capture the seasonal interaction effects.
Wind speed and wind power show a huge amount of autocorrelation, as shown by Ambach and Schmid (2015). Hence, we consider a multivariate seasonal VARMA class model to capture the persistence as well as the periodicity. A VARMA model is also used by Erdem and Shi (2011) to predict a tuple of wind speed time series. In a more general setting, Jeon and Taylor (2012) take a bivariate VARMA generalized autoregressive conditional heteroscedastic (GARCH) approach to model the wind speed and wind direction and convert the predictions of both into wind power forecasts.
Instead of using wind speed as a quadratic regressor as done in the WPPT and GWPPT approach, we use thresholds and vector autoregression to cover the non-linearity. The threshold autoregressive approach is also applied in a context of probabilistic load forecasting by Ziel and Liu (2016) as a suitable tool to explore the non-linearity in the data. Here, we use a VARMA model to capture the correlation structure of several turbines and to predict wind speed and wind power altogether. Finally, we propose a threshold GARCH (TGARCH) model for the wind speed series and a power-TGARCH process for the volatility. With it, we are able to capture the conditionally heteroscedastic behavior in the data, similarly to Ewing et al. (2006) and Ambach and Schmid (2015).
The assumed statistical model structure for the wind speed and power allows us to simulate sample paths for several scenarios. Using bootstrap simulation techniques we can easily derive probabilistic forecasts. As pointed out by, e.g., Pinson et al. (2013), Alessandrini et al. (2013), and Hong et al. (2016), the importance of probabilistic wind power forecasting is increasing, especially for longer forecasting horizons. Zugno et al. (2012) use probabilistic forecasts of wind power as well, Gneiting and Raftery (2007) provide details on the computation and evaluation of probabilistic forecasts. Recently, Berner et al. (2015) discuss bias correction and accuracy improvements in

general probabilistic forecasting. Using a mesoscale meteorology framework, they address the main reason for using probabilistic forecasts instead of point forecasts, i.e. "[... to] account for certain aspects of structural model uncertainty".
For the estimation, we apply a high-dimensional shrinkage technique based on the popular least absolute shrinkage and selection operator (lasso) method, as introduced by Tibshirani (1996). Similarly, Evans et al. (2014) use the lasso method to augment the forecasting accuracy of a wind farm. According to Ziel et al. (2015), we apply an iteratively re-weighted lasso approach to estimate the model parameters. Thus, we can provide a huge parameter space, still come up with a parsimonious and regularized specification and have very convenient computing times in comparison to the usual maximum likelihood technique (i.e. few seconds compared to several minutes on a modern computer). For the time varying and periodic effects, the algorithm will estimate parameters that may vary over time at a certain significance. Otherwise, the parameters remain constant. Our time varying periodic TVARMA-power-TGARCH model returns more accurate forecasts than the usual WPPT and GWPPT models as well as a set of benchmark models, including the usual persistence forecaster. Results show that our model provides less skewed forecast errors than our benchmarks.
This paper makes two major contributions: First, we present a modeling framework for wind power that includes wind speed, flexible modeling of the periodicity and heteroscedasticity. Second, we show how to estimate the model parameters by applying a re-weighted heteroscedastic lasso approach to a time series setting, as has been done recently by Ziel (2015). Empirical results from out-of-sample forecasts are compared to a set of benchmark models.
The paper is structured as follows: Section 2 discusses the data set used. In Section 3 we show our new model idea. Section 4 presents the estimation technique. Empirical results are discussed in Section 5 and Section 6 concludes.

2. Data and Their Characteristics

The turbine data set used in this paper is a high-frequency series collected from a wind park in Germany. The wind park consists of 8 turbines. The observed park is situated in a mostly plain and rural region. The area has a slight roughness with fields and some forestation. Due to a non-disclosure

agreement, the specific locations cannot be revealed. However, Figure 1 presents a stylized map of the turbines' arrangement. The turbines, labeled Turbine A to H, exhibit a power range of $[0; 1500]$ kW each and write sensor data to log files at a frequency of ten minutes. The observed time frame spans from November 1, 2010 to November 5, 2012, so there are 105984 observations per turbine.

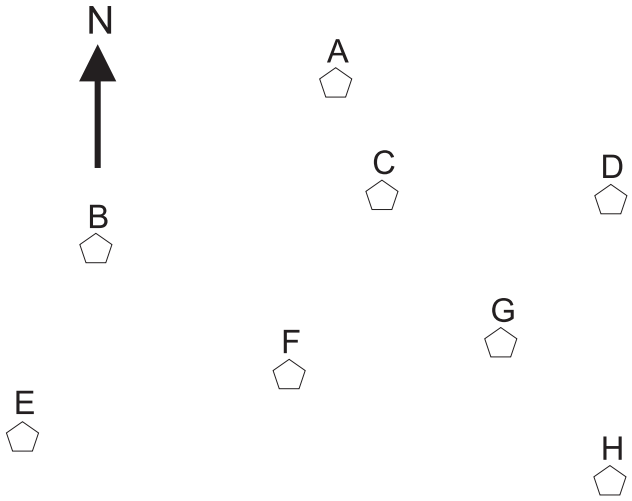

Figure 1: Stylized map of the wind parks investigated.

Table 1 shows descriptive statistics for two of the turbines in the data set. Note that wind power observations may very well be slightly below zero: If wind speeds are below cut-in speed (i.e. there is no hub rotation and thus, no power production), the turbine consumes power for system operation and aviation lights. Also, nacelle and rotor pitch adjustments require appreciable amounts of electricity and are mostly performed during times at which the turbines do not produce power themselves. Thus, some power observations are below zero. Considering the entire data range (from -19 to 1542), the part of values below zero is only about 1% of the range. Thus, we do not consider turbine power consumption to be of much importance. The histograms in Figure 2 support that determination. This also holds true for the very few observations at which the theoretical power maximum of 1500 kW is exceeded. Hence, we stay with the theoretical range assumption of $[0; 1500]$ kW.
The data set possesses a minor number of missing values due to engine error, maintenance shutdown or ice error. About 3% of the data are missing, but

the gaps are small: The maximum run length of missing values is 586, which is about 0.007% of the entire data set. Therefore, we easily fill the gaps by simple linear interpolation.
A scatter plot (empirical power curve) and time series plots of wind speed and wind power are given in Figure 2. It shows that wind speed and wind power follow a similar structure and are closely interdependent.

Statistic	Min	Median	Max	Mean	SD
Speed A	0.4	5.2	18.0	5.1	2.4
Power A	-19.0	150.0	1532.0	217.2	272.0
Speed B	0.4	5.5	18.6	5.3	2.5
Power B	-19.0	155.0	1493.0	230.5	291.1

Table 1: Descriptive statistics of Turbines A and B. Wind speed denoted in m/s, wind power in kW.

Besides the high persistence of wind speed and wind power which is directly related to the high-frequency data set, it is necessary to discuss another important characteristic of wind power and wind speed. The wind speed data provides a strong periodic behavior, as Zhu et al. (2014) and Ambach and Schmid (2015) point out. The wind power data set also provides these characteristics. A diurnal periodicity as considered for the WPPT and GWPPT is observable for our data set, but the annual period is not completely straight forward. Hence, we calculate the sample periodogram, which is shown in Figure 3.
Figure 3 shows the estimated spectral density for the wind power and wind speed of turbine A. The red lines in Figure 3 show annual and half-annual frequencies in the upper panels and daily and half-daily periods in the lower panels. The diurnal periodicity is not so prevalent within the wind power series shown in the right-hand panels, but there are several important frequencies nearby the daily period. After all, periodic B-spline functions will help to model all multiples of a diurnal and annual periodicity.

3. Model

Let $d = 8$ denote the number of turbines in the wind park. Thus, the set of turbines is $\mathcal{D} = \{1, \dots, d\}$. The d-dimensional time series of wind speed

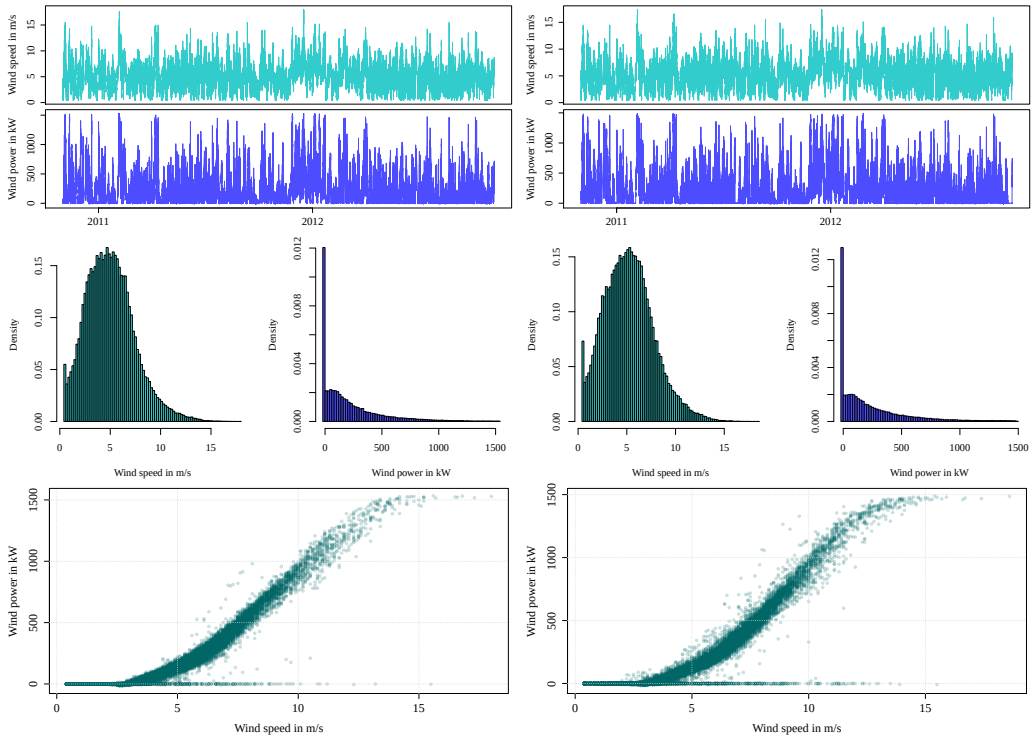

(a) Turbine A. (b) Turbine B.

Figure 2: Time series, histograms of wind speed and power and corresponding empirical power curves of Turbines A and B.

is $(\boldsymbol{W}_t)_{t\in\mathbb{Z}}$ with $\boldsymbol{W}_t = (W_{1,t}, \ldots, W_{d,t})'$ and the wind power is $(\boldsymbol{P}_t)_{t\in\mathbb{Z}}$ with $\boldsymbol{P}_t = (P_{1,t}, \ldots, P_{d,t})'$.
We split the model description into two parts: First, we present the multivariate time varying threshold VARMA model for the wind speed. Its heteroscedastic variance structure is modeled by a TGARCH type process. Afterward, we present the wind power model which includes the wind speed dependence and considers the errors of wind power themselves to follow a power-TGARCH process.

3.1. The Wind Speed Component

For the wind speed $\boldsymbol{W}_t$ we consider the multivariate time varying threshold-VARMA model

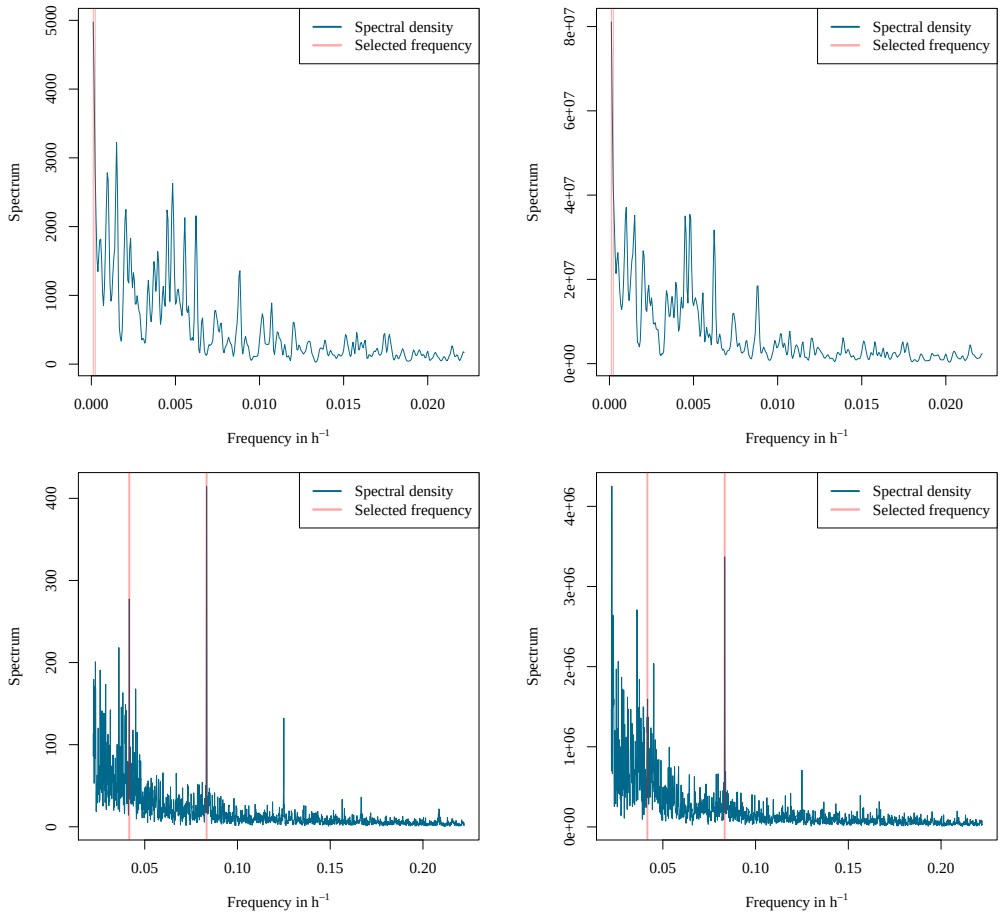

(a) Smoothed Periodogram of wind speed for Turbine A.

(b) Smoothed Periodogram of wind power for Turbine A.

Figure 3: Estimated spectral density of wind speed (left panels) and wind power (right panels).

$$W_{i,t} = \phi_{i,0}(t) + \sum_{j \in \mathcal{D}} \sum_{k \in I^{\phi}_{i,j}} \sum_{c \in C^{\phi}_{i,j,k}} \phi_{i,j,k,c}(t) \max\{W_{j,t-k}, c\} + \sum_{j \in \mathcal{D}} \sum_{k \in I^{\theta}_{i,j}} \theta_{i,j,k}(t) \varepsilon_{j,t-k} + \varepsilon_{i,t}, \tag{1}$$

where $i \in \mathcal{D}$, $\phi_{i,j,k,c}$ resp. $\theta_{i,j,k}$ represent the time varying autoregressive

and moving average coefficients and $\varepsilon_{i,t}$ is the error term. The index sets $I^{\phi}_{i,j}$ and $I^{\theta}_{i,j}$ contain the corresponding relevant AR- and MA-lags and the threshold set $C^{\phi}_{i,j,k}$ contains all considered thresholds in the autoregressive part. The simple choice $C^{\phi}_{i,j,k} = \{-\infty\}$ would turn the model into a standard time varying VARMA process. The thresholds describe the AR-dependence of wind speed by a piecewise linear function with breaks at the corresponding thresholds. Just as each smooth function, it can be approximated well by piecewise linear functions, which provides a flexible and efficient way to capture the non-linear dependence in the data.
We assume the error process $(\varepsilon_{i,t})_{t\in\mathbb{Z}}$ to be conditionally heteroscedastic. Therefore, we consider $\varepsilon_{i,t} = \sigma_{i,t}Z_{i,t}$, where $(Z_{i,t})_{t\in\mathbb{Z}}$ is i.i.d. with $\mathbb{E}(Z_{i,t}) = 0$ and $\mathbb{V}\mathrm{ar}(Z_{i,t}) = 1$. In detail, we assume that $\varepsilon_{i,t}$ follows a time varying TGARCH process, such that

$$\begin{aligned}\sigma_{i,t} = \alpha_{i,0}(t) &+ \sum_{j\in\mathcal{D}}\sum_{k\in I^{\alpha}_{i,j}} \alpha^{+}_{i,j,k}(t)\varepsilon^{+}_{j,t-k} + \alpha^{-}_{i,j,k}(t)\varepsilon^{-}_{j,t-k} \\ &+ \sum_{j\in\mathcal{D}}\sum_{k\in I^{\beta}_{i,j}} \beta_{i,j,k}(t)\sigma_{j,t-k} \qquad (2)\end{aligned}$$

with index sets $I^{\alpha}_{i,j}$ and $I^{\beta}_{i,j}$, $\varepsilon^{+}_{j,t-k} = \max\{\varepsilon_{j,t-k}, 0\}$, $\varepsilon^{-}_{j,t-k} = \max\{-\varepsilon_{j,t-k}, 0\}$ and time varying coefficients $\alpha_{i,0}(t) > 0$, $\alpha^{+}_{i,j,k}(t) \geq 0$, $\alpha^{-}_{i,j,k}(t) \geq 0$, $\beta_{i,j,k}(t) \geq 0$. The index sets for the considered lags are given in Table 2.
For most of the coefficients, we allow dependence of up to one hour (6 lags) only to keep the specification manageable. However, for the coefficients that describe the wind dependence on its own past, we allow for more parameters. Here, we also include the lags $140, \ldots, 150$ to cover the impact from the previous day (which corresponds to $6 \times 24 = 144$ lags).
For the thresholds $C^{\phi}_{i,j,k}$, we use a parsimonious lag specification: We allow non-linear impacts for the first two lags, only. The elements of $C^{\phi}_{i,j,k}$ contain the 10% percentiles of the process in the mean equation. Thus, $C^{\phi}_{i,j,k}$ contains the 10% percentiles in the cases $k = 1$ or $k = 2$. All elements that do not satisfy this restriction are set to $C^{\phi}_{i,j,k} = \{-\infty\}$. The non-linear impact of the threshold model specification acts via piecewise linear functions to cover possibly present turbulent flow and wake effects. The effect of this model component is explained in detail in the following subsection.
To keep the parameter space reasonable, we keep most of the coefficients

constant and allow only a few important ones to vary over time. For $I^{\phi}_{i,j}$ as well as $I^{\theta}_{i,i}$, we consider the coefficients for lags 1 and 2 to be time varying. Ziel et al. (2015) proceed similarly for modeling the wind and solar power net feed-in.

Index sets	Contained lags
$I^{\phi}_{i,i}$, $I^{\alpha}_{i,i}$	1, ..., 40 and 140, ..., 150
$I^{\phi}_{i,j}$, $I^{\theta}_{i,i}$, $I^{\theta}_{i,j}$, $I^{\alpha}_{i,j}$,$I^{\beta}_{i,i}$, $I^{\beta}_{i,j}$	1, ..., 6

Table 2: Considered lags of the index sets, where $i, j \in \mathcal{D}$ with $j \neq i$.

3.2. The Wind Power Model

For the wind power process we assume a model that is given by

$$\begin{aligned} P_{i,t} =& \varphi_{i,0}(t) + \sum_{j\in\mathcal{D}} \sum_{k\in I^{\varphi}_{i,j}} \sum_{c\in C^{\varphi}_{i,j,k}} \varphi_{i,j,k,c}(t) \max\{P_{j,t-k}, c\} \\ &+ \sum_{j\in\mathcal{D}} \sum_{k\in I^{\psi}_{i,j}} \sum_{c\in C^{\psi}_{i,j,k}} \psi_{i,j,k,c}(t) \max\{W_{j,t-k}, c\} \\ &+ \sum_{j\in\mathcal{D}} \sum_{k\in I^{\vartheta}_{i,j}} \vartheta_{i,j,k}(t)\epsilon_{j,t-k} \\ &+ \sum_{j\in\mathcal{D}} \sum_{k\in I^{\varpi}_{i,j}} \varpi_{i,j,k}(t)\varepsilon_{j,t-k} + \epsilon_{i,t}. \end{aligned} \tag{3}$$

All observed turbines are located in close proximity to each other, and are influenced by the same air pressure and weather conditions. Dependent on the angle of movement of a particular pressure area (and thus, wind conditions) at any one time, these conditions may hit one set of turbines sooner than others. The spatial dispersion of the turbines' power production can therefore be accounted for by a time-lag structure. Thus, we assume that the power P_i of turbine i can depend on its own past as well as on the past of the power of the other turbines by the φ parameters. The power can also depend on the current and past wind speed by means of the ψ parameters. Furthermore, wind power depends on the past residuals. Note that the lag structure for the wind power model is slightly different from that of the wind

speed model, as we assume a causal/temporal structure in the data. We assume that the wind speed $\boldsymbol{W}_t$ at time t can only depend on the past wind speed $\boldsymbol{W}_{t-k}$ for $k \geq 1$. Similarly, the wind power $\boldsymbol{P}_t$ depends on the past wind power $\boldsymbol{P}_{t-k}$ for $k \geq 1$, but also on current and past wind speed $\boldsymbol{W}_{t-k}$ for $k \geq 0$.
Comparably to the wind speed model, we allow for non-linear wind speed and wind power effects by several thresholds. Particularly, the theoretical non-linear effect of the wind speed on the wind power is well known to be described by the third-degree polynomial:

$$P = \frac{1}{2}\rho C_P A W^3, \tag{4}$$

where ρ describes the air density, C_P denotes the physical properties of the turbine (values of up to 16/27, the so-called Betz limit), and A represents the swept area. Hennessey (1977) goes into details. However, especially around the upper bound of the maximum produced wind power, it is known that the true impact of wind speed on wind power is different from the usual cubic relationship and should not be modeled to be cubic. Our way of modeling the non-linear impact by piecewise linear effects allows for a flexible way to model the underlying non-linear impact. To illustrate the impact of the thresholds we briefly present a simple threshold model for the wind power dependent on the wind speed. It is given by

$$P_{i,t} = a + \sum_{c=0}^{16} b_c \max\{W_{i,t}, c\} + \mathrm{e}_{i,t} \tag{5}$$

for turbine i. It contains thresholds at $0, 1, 2, \ldots, 16$ for modeling the non-linear relationship by piecewise linear functions. In Figure 4, the fitted values of model (5) are given for Turbines A and B of the investigated wind park. It can be seen that the piecewise linear approach is able to cover the non-linear relationship quite well. There are distinct bents at the threshold points. In fact, the fitted curve is a linear spline. Of course, a higher number of thresholds will increase the model fit. However, if the number of thresholds is too large, it might lead to overfitting. Still, this problem is somehow limited due to our shrinkage estimation procedure, so that a large number of parameters in the problem space does not necessarily imply a lot of estimations in the solution. Instead, the algorithm will automatically select the

most plausible piecewise linear function that approximates the non-linear impact well, as we choose the thresholds in model (3) to be data driven.
Similarly as for the wind speed process, we assume a GARCH-type process for the wind power error, so $\epsilon_{i,t} = \varsigma_{i,t}U_{i,t}$ with $U_{i,t}$ i.i.d., $\mathbb{E}(U_{i,t}) = 0$ and $\mathbb{E}(U_{i,t}^2) = 1$. This slightly differs from the TGARCH process for the wind speed: We assume that the third-degree relationship in equation (4) hands down to the residual volatility. Thus, instead of considering a recursion on $\varsigma_{i,t}$, we consider a recursion on the cubed root of the volatility $\varsigma_{i,t}^{\frac{1}{3}}$. Consequently, we assume that $\epsilon_{i,t}$ follows a time varying power-TGARCH process:

$$\begin{aligned}\varsigma_{i,t}^{\frac{1}{3}} =& \eta_{i,0}(t) + \sum_{j\in\mathcal{D}}\sum_{k\in I_{i,j}^{\eta}} \eta_{i,j,k}^{+}(t)|\epsilon_{j,t-k}^{+}|^{\frac{1}{3}} + \eta_{i,j,k}^{-}(t)|\epsilon_{j,t-k}^{-}|^{\frac{1}{3}} \\ &+ \sum_{j\in\mathcal{D}}\sum_{k\in I_{i,j}^{\zeta}} \zeta_{i,j,k}(t)\varsigma_{j,t-k}^{\frac{1}{3}} + \sum_{j\in\mathcal{D}}\sum_{k\in I_{i,j}^{\upsilon}} \upsilon_{i,j,k}^{+}(t)|\varepsilon_{j,t-k}^{+}|^{\frac{1}{3}} \\ &+ \upsilon_{i,j,k}^{-}(t)|\varepsilon_{j,t-k}^{-}|^{\frac{1}{3}} + \sum_{j\in\mathcal{D}}\sum_{k\in I_{i,j}^{\varrho}} \varrho_{i,j,k}(t)\sigma_{j,t-k}^{\frac{1}{3}}, \end{aligned} \tag{6}$$

$\epsilon_{j,t-k}^{+} = \max\{\epsilon_{j,t-k}, 0\}$, $\epsilon_{j,t-k}^{-} = \max\{-\epsilon_{j,t-k}, 0\}$. Finally, the index sets for the considered lags on the mean part of the model in equation (3) and the variance part in equation (6) are given in Table 3. The corresponding parameters for the index sets $I_{i,j}^{\alpha}$, $I_{i,j}^{\beta}$, $I_{i,j}^{\varphi}$, $I_{i,j}^{\vartheta}$, $I_{i,j}^{\eta}$, $I_{i,j}^{\upsilon}$, $I_{i,j}^{\zeta}$ and $I_{i,j}^{\varrho}$ are considered to be time varying on lags 1 and 2, those for the sets $I_{i,j}^{\psi}$ and $I_{i,j}^{\varpi}$ are time varying on lags $0, 1, 2$ and the corresponding regressors are modeled by periodic B-splines, as discussed subsequently.

3.3. Time Varying Coefficients

We assume an identical structure of the time varying coefficients in the wind speed and in the wind power model, as both exhibit similar seasonal effects. In general, the time varying coefficients can be modeled by periodic functions like Fourier approximations or other periodic basis functions, such as periodic B-splines or periodic wavelets. We opt for the flexible cubic B-spline approach. Let ξ be a time varying coefficient. Then

$$\xi(t) = \sum_{l=1}^{N_{\xi}} \xi_l B_l^{\xi}(t), \tag{7}$$

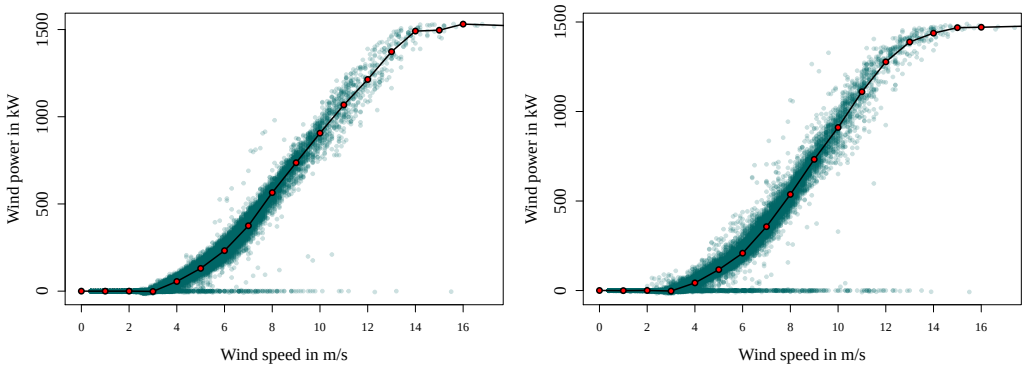

(a) Turbine A. (b) Turbine B.

Figure 4: Fitted results of the illustrative example of model (5).

Index sets	Contained lags
$I^{\varphi}_{i,i}$, $I^{\eta}_{i,i}$, $I^{\upsilon}_{i,i}$	$1, \ldots, 40$ and $140, \ldots, 150$
$I^{\psi}_{i,i}$	$0, \ldots, 40$ and $140, \ldots, 150$
$I^{\vartheta}_{i,i}$, $I^{\varphi}_{i,j}$, $I^{\vartheta}_{i,j}$, $I^{\eta}_{i,j}$, $I^{\upsilon}_{i,j}$, $I^{\zeta}_{i,i}$, $I^{\zeta}_{i,j}$, $I^{\varrho}_{i,i}$, $I^{\varrho}_{i,j}$	$1, \ldots, 6$
$I^{\varpi}_{i,i}$, $I^{\psi}_{i,j}$, $I^{\varpi}_{i,j}$	$0, \ldots, 6$

Table 3: Considered lags of the index sets, where $i, j \in \mathcal{D}$ with $j \neq i$.

so that ξ is given by a sum of N_ξ basis functions $B^{\xi}_{l}(t)$, weighted by ξ_l.[1] In the literature on wind power modeling, Fourier approximations are used frequently, see, e.g., Giebel et al. (2011). However, the Fourier technique is a global approach. For our purpose, a local approach is preferable, as it is more flexible with respect to possible changes in the time-dependent structure itself. We design our basis function so that it can cover both the diurnal and the annual periodic effects. Furthermore, we allow for possible interactions between both seasonalities, so that the diurnal impact can change over the

[1]Details on the construction of the B-splines basis functions are discussed in the appendix.

year. This impact is visualized in Figure 5. It displays the daily mean wind speed and wind power of all considered wind turbines for the four seasons in a year. For both the wind speed and the wind power, we observe that during the morning hours around 7am to 10am, there is a distinct drop, which is less severe during the winter months. Moreover, it can be seen that in the summer, this drop happens earlier than in the other seasons (from around 6am to 8am). In contrast, during winter time, this drop seems to happen quite late in the day (from around 9am to 10am), best visible in Figure 5a. This indicates strong interaction of the wind speed with the sunrise. Over all, the daily mean curves differ significantly from each other, showing that the diurnal pattern depends on the annual pattern, and vice versa.

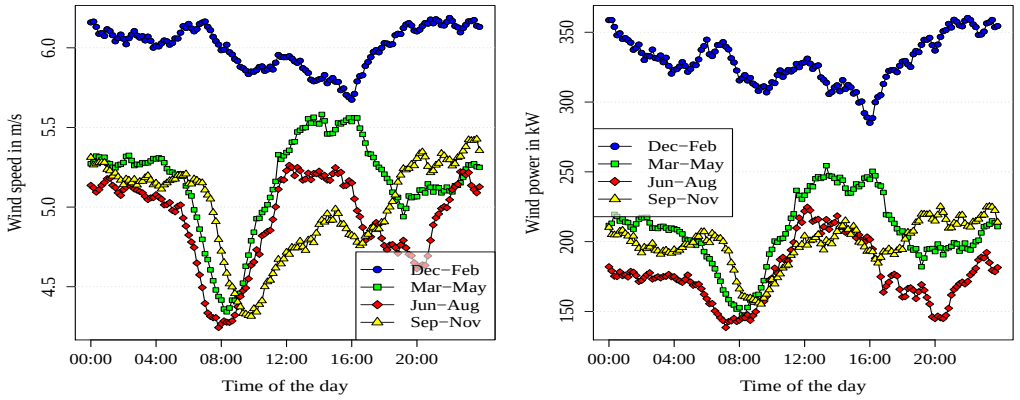

(a) Mean wind speed.

(b) Mean wind power.

Figure 5: Daily mean wind speed and wind power of the 8 turbines for the four seasons in a year.

4. Estimation

As an estimation algorithm, we use a lasso based estimation technique. Lasso is a penalized least square regression method. Thus, we consider the least squares representation of our model. For the conditional mean models (1) and (3), the regression representations are given by

$$\mathcal{W}_i = \mathbb{W}_i \boldsymbol{b}_i^{\mathcal{W}} + \mathcal{E}_i \tag{8}$$

$$\mathcal{P}_i = \mathbb{P}_i \boldsymbol{b}_i^{\mathcal{P}} + E_i. \tag{9}$$

Here, $\mathcal{W}_i = (W_{i,1}, \ldots, W_{i,n})'$ and $\mathcal{P}_i = (P_{i,1}, \ldots, P_{i,n})'$ denote the observed wind speed and wind power vectors, $\mathbb{W}_i$ and $\mathbb{P}_i$ are the matrices of covariates that correspond to (1) and (3), $\boldsymbol{b}_i^{\mathcal{W}}$ and $\boldsymbol{b}_i^{\mathcal{P}}$ are the full parameter vectors, $\mathcal{E}_i = (\varepsilon_{i,1}, \ldots, \varepsilon_{i,n})'$ and $E_i = (\epsilon_{i,1}, \ldots, \epsilon_{i,n})'$ denote the error vectors and n is the number of observations.
Similarly, we formulate a regression representation for the volatility models. We exploit the fact that $|\varepsilon_{i,t}| = \gamma_i \sigma_{i,t} + \sigma_{i,t}(|Z_{i,t}| - \gamma_i)$ and $|\epsilon_{i,t}|^{\frac{1}{3}} = \tau_i \varsigma_{i,t}^{\frac{1}{3}} + \varsigma_{i,t}^{\frac{1}{3}}(|U_{i,t}|^{\frac{1}{3}} - \tau_i)$, where $\gamma_i = \mathbb{E}|Z_{i,t}|$ and $\tau_i = \mathbb{E}|Z_{i,t}|^{\frac{1}{3}}$. Note that $v_{i,t} = \sigma_{i,t}(|Z_{i,t}| - \gamma_i)$ and $u_{i,t} = \varsigma_{i,t}(|U_{i,t}| - \tau_i)$ are weak white noise processes. With that we express a recursion on $|\varepsilon_{i,t}|$ and $|\epsilon_{i,t}|^{\frac{1}{3}}$ by

$$\begin{aligned}
|\varepsilon_{i,t}| =& \gamma_i \alpha_{i,0}(t) + \sum_{j \in \mathcal{D}} \sum_{k \in I_{i,j}^{\alpha}} \gamma_i \alpha_{i,j,k}^{+}(t) \varepsilon_{j,t-k}^{+} + \gamma_i \alpha_{i,j,k}^{-}(t) \varepsilon_{j,t-k}^{-} \\
&+ \sum_{j \in \mathcal{D}} \sum_{k \in I_{i,j}^{\beta}} \gamma_i \beta_{i,j,k}(t) \sigma_{j,t-k} + v_{i,t} \qquad (10) \\
|\epsilon_{i,t}|^{\frac{1}{3}} =& \tau_i \eta_{i,0}(t) + \sum_{j \in \mathcal{D}} \sum_{k \in I_{i,j}^{\eta}} \tau_i \eta_{i,j,k}^{+}(t) |\epsilon_{j,t-k}^{+}|^{\frac{1}{3}} + \eta_{i,j,k}^{-}(t) |\epsilon_{j,t-k}^{-}|^{\frac{1}{3}} \\
&+ \sum_{j \in \mathcal{D}} \sum_{k \in I_{i,j}^{\zeta}} \tau_i \zeta_{i,j,k}(t) \varsigma_{j,t-k}^{\frac{1}{3}} \\
&+ \sum_{j \in \mathcal{D}} \sum_{k \in I_{i,j}^{\upsilon}} \tau_i \upsilon_{i,j,k}^{+}(t) |\varepsilon_{j,t-k}^{+}|^{\frac{1}{3}} + \upsilon_{i,j,k}^{-}(t) |\varepsilon_{j,t-k}^{-}|^{\frac{1}{3}} + \\
&+ \sum_{j \in \mathcal{D}} \sum_{k \in I_{i,j}^{\varrho}} \tau_i \varrho_{i,j,k}(t) \sigma_{j,t-k}^{\frac{1}{3}} + u_{i,t}. \qquad (11)
\end{aligned}$$

The corresponding multivariate regression representations are given by

$$|\mathcal{E}_i| = \mathbb{V}_i \boldsymbol{a}_i^{\mathcal{E}} + \mathcal{V}_i \qquad (12)$$
$$|E_i| = \mathbb{U}_i \boldsymbol{a}_i^{E} + \mathcal{U}_i, \qquad (13)$$

where $\mathbb{V}_i$ and $\mathbb{U}_i$ are the regressor matrices that correspond to (10) and (11), $\boldsymbol{a}_i^{\mathcal{E}}$ and $\boldsymbol{a}_i^{E}$ are the parameter vectors and $\mathcal{V}_i = (v_{i,1}, \ldots, v_{i,n})$ and $\mathcal{U}_i = (u_{i,1}, \ldots, u_{i,n})$.
For the parameter estimation we use an estimation technique that is based

on a lasso regression for heteroscedastic data. The approach is similar to the popular FGLS (feasible generalized least squares) solution by Newey and West (1987) at which a weighting matrix brings the "meat" into the estimation, leading to point-wise heteroscedasticity consistency. This lasso method was analyzed first by Wagener and Dette (2012) in a standard regression setting and by Ziel (2015) in a time series setting. We slightly modify the algorithm to plug-in the causal setting, i.e. the multivariate approach for the wind power mean model. Therefore, we apply a weighted lasso for the conditional mean regressions (8) and (9). For the conditional variance regressions (12) and (13), we just apply a standard lasso.
Let $\boldsymbol{\Omega}_i = \text{diag}(\boldsymbol{\omega}_i)$ and $\boldsymbol{\Xi}_i = \text{diag}(\boldsymbol{\xi}_i)$ be diagonal matrices of heteroscedasticity weights $\boldsymbol{\omega}_i = (\omega_{i,1}, \dots, \omega_{i,n})$ and $\boldsymbol{\xi}_i = (\xi_{i,1}, \dots, \xi_{i,n})$. The weighted lasso optimization problems concerning (8), (9), (12) and (13) are given by

$$\widehat{\boldsymbol{b}}_i^{\mathcal{W}} = \arg\min_{\boldsymbol{b}} (\mathcal{W}_i - \mathbb{W}_i \boldsymbol{b})' \boldsymbol{\Omega} (\mathcal{W}_i - \mathbb{W}_i \boldsymbol{b}) + \lambda_i^{\mathcal{W}} |\boldsymbol{b}| \tag{14}$$

$$\widehat{\boldsymbol{b}}_i^{\mathcal{P}} = \arg\min_{\boldsymbol{b}} (\mathcal{P}_i - \mathbb{P}_i \boldsymbol{b})' \boldsymbol{\Xi} (\mathcal{P}_i - \mathbb{P}_i \boldsymbol{b}) + \lambda_i^{\mathcal{P}} |\boldsymbol{b}| \tag{15}$$

$$\widehat{\boldsymbol{a}}_i^{\mathcal{E}} = \arg\min_{\boldsymbol{a} \geq \boldsymbol{0}} (|\mathcal{E}_i| - \mathbb{U}_i \boldsymbol{a})' (|\mathcal{E}_i| - \mathbb{U}_i \boldsymbol{a}) + \lambda_i^{\mathcal{E}} |\boldsymbol{a}| \tag{16}$$

$$\widehat{\boldsymbol{a}}_i^{E} = \arg\min_{\boldsymbol{a} \geq \boldsymbol{0}} (|E_i| - \mathbb{V}_i \boldsymbol{a})' (|E_i| - \mathbb{V}_i \boldsymbol{a}) + \lambda_i^{E} |\boldsymbol{a}|, \tag{17}$$

with tuning parameters $\lambda_i^{\mathcal{W}}$, $\lambda_i^{\mathcal{P}}$, $\lambda_i^{\mathcal{E}}$ and λ_i^{E}. Here, we only allow estimators $\widehat{\boldsymbol{a}}_i^{\mathcal{E}}$ and $\widehat{\boldsymbol{a}}_i^{E}$ with no negative entry to ensure that the recurrence equation (and thus, the volatilities $\sigma_{i,t}$ and $\varsigma_{i,t}$) are well defined.
For solving the lasso optimization problems we use the coordinate descent algorithm as introduced by Friedman et al. (2007). This algorithm solves the problem on a given tuning parameter grid. We choose the tuning parameters by the minimalist but rather strict Bayesian information criterion (BIC), to avoid overfitting. Other information criteria or a cross-validation based approach can be also be applied.
In the first step, we estimate the conditional mean parameters $\boldsymbol{b}^{\mathcal{W}}$ and $\boldsymbol{b}^{\mathcal{P}}$. Then, we consider the estimated residuals for the estimation of the volatility parameters $\boldsymbol{a}^{\mathcal{E}}$ and $\boldsymbol{a}^{E}$. Afterward, we use the fitted volatilities to redefine the heteroscedasticity matrices $\boldsymbol{\Omega}_i$ and $\boldsymbol{\Xi}_i$ to repeat the procedure with the new weight matrices for the conditional mean parameters $\boldsymbol{b}^{\mathcal{W}}$ and $\boldsymbol{b}^{\mathcal{P}}$. In practice however, there is an initialization problem, as the residuals $\varepsilon_{i,t}$ and

$\epsilon_{i,t}$ as well as the conditional standard deviations $\sigma_{i,t}$ and $\varsigma_{i,t}$ are unknown. Still, Chen and Chan (2011) discuss a method for estimating ARMA models using an iterative lasso approach. Their algorithm basically uses the fact that every ARMA(p, q) can be written as an AR(∞). Thus, it can be approximated by an AR(p) for large p. This is also applicable for time varying threshold VARMA models. Thus, a time varying threshold VARMA(p, q) process is a time varying threshold AR(∞) process. Similarly, for the volatility model, it holds that every exponentiated ARMA model can be expressed as a PGARCH with the corresponding power. For a power of two, a squared ARMA process is a GARCH process. So, the same relationship holds: A time varying power-GARCH(p, q) process is a time varying power-ARCH(∞) process.
Using these facts we can handle the initialization problem as follows: In the first iteration step, we replace all $\varepsilon_{i,t}$, $\epsilon_{i,t}$, $\sigma_{i,t}$ and $\varsigma_{i,t}$ by 1. Thus, in the first step, we actually estimate time varying threshold AR processes in the conditional mean equations and time varying power-ARCH processes in the second step.
Finally, we initialize the heteroscedasticity weights $\boldsymbol{\Omega}_i$ and $\boldsymbol{\Xi}_i$. Here, we simply assume homoscedasticity in the first step, so we take $\boldsymbol{\Omega}_i = \boldsymbol{\Xi}_i = I$. The algorithm can be stated as follows:

1) Initialize $\mathbb{W}_i$, $\mathbb{P}_i$, $\boldsymbol{\Omega}_i$, $\boldsymbol{\Xi}_i$ for all $i \in \mathcal{D}$ and $K = 1$.
2) Estimate $\widehat{\boldsymbol{b}}_i^{\mathcal{W}}$ and $\widehat{\boldsymbol{b}}_i^{\mathcal{P}}$ with (14) and (15) using coordinate descent with weights $\boldsymbol{\Omega}_i$ and $\boldsymbol{\Xi}_i$ for all $i \in \mathcal{D}$.
3) Estimate $\widehat{\boldsymbol{a}}_i^{\mathcal{E}}$ and $\widehat{\boldsymbol{a}}_i^{E}$ with (16) and (17) by coordinate descent using the estimated residuals $(\widehat{\varepsilon}_{i,1}, \ldots, \widehat{\varepsilon}_{i,n})$ and $(\widehat{\epsilon}_{i,1}, \ldots, \widehat{\epsilon}_{i,n})$ from *2)* to compute $\mathbb{U}_i$ and $\mathbb{V}_i$ for all $i \in \mathcal{D}$.
4) Compute the estimated volatilities $(\widehat{\sigma}_{i,1}, \ldots, \widehat{\sigma}_{i,n})$ and $(\widehat{\varsigma}_{i,1}, \ldots, \widehat{\varsigma}_{i,n})$ with the fitted values from *3)* and redefine $\mathbb{W}_i$, $\mathbb{P}_i$, $\boldsymbol{\Omega}_i = \operatorname{diag}(\boldsymbol{\omega}_i)$ and $\boldsymbol{\Xi}_i = \operatorname{diag}(\boldsymbol{\xi}_i)$ by $\boldsymbol{\omega}_i = (\widehat{\sigma}_{i,1}^{-2}, \ldots, \widehat{\sigma}_{i,n}^{-2})$ and $\boldsymbol{\xi}_i = (\widehat{\varsigma}_{i,1}^{-2}, \ldots, \widehat{\varsigma}_{i,n}^{-2})$ for all $i \in \mathcal{D}$.
5) If $K < K_{\max}$ then $K = K + 1$ and back to *2)*, otherwise stop the algorithm.

We stop the algorithm after a maximum of $K_{\max} = 2$ iterations, which already provides a good ratio of accuracy and computing time. Ziel (2015)

shows that under some regularity conditions, two iterations are sufficient to receive optimal asymptotic properties. As the considered estimation methodology is based on the coordinate descent algorithm, it shares the same computational complexity. With n as number of observations, d as number of turbines and p as dimension of the underlying lasso problem (dimension of parameter vector in (9)), the asymptotic computational complexity of the algorithm is $\mathcal{O}(dnp)$. Thus, if either d, n or p is doubled, the computation time gets doubled as well. Notably, the estimation procedure is easily applicable for large wind parks.

5. Forecasting and Results

After the estimation, the obtained parameters are fit to the current set of data in order to calculate a forecast. As it is the very nature of the lasso approach to return a lot of zero valued parameters, the high-dimensional parameter space is shrunk to a manageable amount of relevant parameters, conditional on the unique settings of in-sample data at each point forecast. We evaluate our model ("lasso") and several benchmark approaches according to their forecasting accuracy. The common criterion is the mean absolute error (MAE). The out-of-sample (OOS) forecasts are performed for a time frame from November 2011 to November 2012. Most benchmark models require appreciable amounts of computing time (several minutes per forecast). To keep the time consumption reasonable, we select $N = 1000$ points in time $(\chi^{(l)}, l = 1, \ldots, N)$ in the out-of-sample period at random. For the respective in-sample periods, we consider the corresponding preceding year with 52830 observations each. Forecasts are calculated at horizons of up to a maximum of two days (i.e. 48 hours = 288 steps). MAE is calculated by

$$\mathrm{MAE}_{i,k} = \frac{1}{N} \sum_{l=1}^{N} \left| P_{i,\chi^{(l)}+k} - \widehat{P}_{i,\tau^{(l)}+k} \right|, \tag{18}$$

where $\widehat{P}_{i,\chi^{(l)}+k}$ is the k-step forecast of wind power and $P_{i,\chi^{(l)}+k}$ is the corresponding actual observation, each at station i. As the results look similar for all of the eight turbines, we just report the mean results over all turbines, so we evaluate

$$\mathrm{MAE}_k = \frac{1}{d}\sum_{i=1}^{d}\mathrm{MAE}_{i,k}. \tag{19}$$

Results for the distinct turbines are available upon request. Additionally to the MAE_k we compute the difference of MAE_k to the persistent benchmark model, denoted by DMAE_k, i.e.

$$\mathrm{DMAE}_k = \mathrm{MAE}_k - \mathrm{MAE}_k^{\mathrm{pers.}}, \tag{20}$$

where $\mathrm{MAE}_k^{\mathrm{pers.}}$ is the MAE_k of the persistent forecaster. The persistent forecaster (so-called naïve predictor, $\hat{P}_{\chi^{(l)}+k} = P_{\chi^{(l)}}$, see, e.g., Costa et al., 2008) is suitable to illustrate the improvement of sophisticated forecasting models in direct comparison to this common quasi-standard benchmark.
We compare our model's results to further benchmarks. We consider a simple univariate AR on the wind power on each turbine $i \in \mathcal{D}$ (AR), a bivariate VAR on wind power and wind speed of each turbine $i \in \mathcal{D}$ (BVAR), a $2 \times d = 16-$dimensional multivariate VAR, jointly on all wind power and wind speed processes (abbr.: VAR), the established WPPT model and its recent generalization, GWPPT. Furthermore, we evaluate an ARMA model, an artificial neural network based approach (ANN), and a gradient boosting machine (GBM).
The AR-type models (AR, BVAR, VAR) are estimated by solving the system of Yule-Walker equations, which guarantees a stationary solution. The corresponding autoregressive order is chosen by minimizing the Akaike information criterion (AIC). Next to AR-type models we consider a univariate ARMA(1,1) process benchmark model for each turbine. De Giorgi et al. (2011) states that out of all ARMA-type models, the ARMA(1,1) process yield the best forecasting results for a short forecasting horizon. We estimate the ARMA model by maximizing the Gaussian likelihood.
The WPPT is based on a turbine specific dynamic regression approach. It takes wind speed as a regressor and captures diurnal periodicity by a Fourier series of time of day observations to estimate the parameters of the model

$$\hat{P}_{t+k} = m + a_1 \cdot P_t + a_2 \cdot P_{t-1} + b_1 \cdot W_{t+k|t} + b_2 \cdot W^2_{t+k|t} + d^c_1 \cdot \cos\left(\frac{2\pi d_{t+k}}{144}\right)$$
$$+ d^c_2 \cdot \cos\left(\frac{4\pi d_{t+k}}{144}\right) + d^s_1 \cdot \sin\left(\frac{2\pi d_{t+k}}{144}\right) + d^s_2 \cdot \sin\left(\frac{4\pi d_{t+k}}{144}\right) + \varepsilon_{t+k}, \quad (21)$$

where $W_{t+k|t}$ is wind speed at time $t+k$ given at time t, d_t is time of day for observation t and ε_{t+k} is assumed white noise.
The generalization of WPPT, GWPPT, is modeled as both-sided censored: Each wind turbine is manufactured to operate at a certain range, the so-called power range. GWPPT makes use of this a-priori known information. The model imposes the following structure on wind power:

$$P^*_t = \eta(\mathbf{z}_t) + \varepsilon_t, \quad (22)$$

where $\mathbf{z}_t$ is the vector of explanatory variables, η is a linear function of $\mathbf{z}_t$, and ε_t is an assumed Gaussian error term. To comply with WPPT, GWPPT assumes the structure shown in equation (21), but adds wind direction to the specification. GWPPT imposes a censored data structure, so that

$$P_t = \begin{cases} l, & P^*_t \leq l \\ P^*_t, & P^*_t \in (l, u) \\ u, & P^*_t \geq u, \end{cases} \quad (23)$$

where l and u are the lower and upper censoring points. Parameters are estimated using a generalized Tobit model. In the end, due to assumed Gaussian errors, the forecast is calculated by

$$\hat{P}_{t+k} = (\Phi(f_2) - \Phi(f_1)) \cdot P^*_{t+k} + (\phi(f_1) - \phi(f_2)) \cdot \widehat{\sigma} + u \cdot (1 - \Phi(f_2)), \quad (24)$$

where $f_1 = (l - P^*_{t+k})/\widehat{\sigma}$, $f_2 = (u - P^*_{t+k})/\widehat{\sigma}$, and $\phi(\cdot)$ and $\Phi(\cdot)$ denote normal PDF (Probability Density Function) and CDF (Cumulative Distribution Function), respectively.
For the ANN benchmark, we consider a single feed-forward neural network and stick close to the setting as used by Li and Shi (2010). As its inputs, we consider the lagged values of the past 8 hours of observations. We train the neural network on 50 training vectors and on 4 neurons in the hidden layer.

The last benchmark investigated is based on gradient boosting machines (GBM). Landry et al. (2016) use GBM methods successfully for the wind power forecasting track in the Global Energy Forecasting Competition 2014. We train a GBM for each wind turbine with a memory of 5 hours on the full data set. Similarly to Landry et al. (2016), we choose the shrinkage tuning parameter to be 0.05, the interaction depth to be 5 and the minimum number of observations to be 30. In total, we choose 100 trees, which are sufficient to reach convergence.

Figure 6 presents the out-of-sample aggregated forecasting error results. Looking at MAE and DMAE, persistence is outperformed by far by most models. ARMA and ANN perform badly. The GBM and the AR type models perform better, but are still not very competitive. Most of the times, lasso competes with WPPT and GWPPT, but sometimes, lasso outruns (G)WPPT, e.g. at forecasting horizons of around 8 hours and above 32 hours.

Table 4 shows the results for several selected forecasting horizons (1 step, 6 steps (1 hour), 24 steps (4 hours), 48 steps (8 hours), 72 steps (12 hours), 144 steps (1 day) and 288 steps (2 days)). As can be seen, lasso is either the best model or not significantly different from the best model.

Note that for longer forecasting horizons (e.g. 24 or 48 hours), the surplus of point forecasting is limited due to the strong amount of uncertainty. However, the proposed model can be used for probabilistic forecasting as well. Using residuals based bootstrap as done by, e.g., Ziel and Liu (2016), we can easily simulate sample paths for the wind speed and power of all turbines in a wind park. We evaluate the empirical quantiles of the bootstrap samples paths and obtain an estimate for the corresponding quantile. Exemplarily, Figure 7 shows the probabilistic wind speed and power forecast for the 99 percentiles for Turbines A and B, starting at February 25th, 2012, 07:20. The figure reveals both, the diurnal seasonal pattern as well as heteroscedasticity. For instance, it can be seen that at a forecasting horizon of 4 as well as for $24 + 4 = 28$ hours, there are greater forecasting values for both the wind speed and the wind power. Indeed, the observations around these peaks are greater than those in the near proximity. Additionally, these peaks are rather volatile, so that the prediction intervals at these peaks are relatively wide, slightly wider than in the neighboring hours. Overall, we see that the prediction intervals get wider with increasing forecasting horizon as expected. In general, they seem to be relatively wide. However, we see that in each of the four figures some observations fall into the reddish colored area

which represents large prediction intervals. Most distinct, in Figure 7d and for large forecasting horizons of more than 40 hours it can be seen that all observations of the wind power of Turbine B fall into the prediction area of very small probabilities. Thus, the prediction intervals do not seem to be too wide or too conservative.

Finally, we investigate the OOS errors' asymmetry by looking at the errors' densities. For more lucidity, we restrain the plots to a few models, lasso, AR and GWPPT. As Figure 8a shows, for the one step ahead forecast, all models return symmetric and leptokurtic results. This symmetry declines for increasing forecasting horizons. For the 24 steps (4 hours) ahead forecast, AR starts to tend to asymmetry, as Figure 8b shows. The average error is negative, which represents a systematic over-estimation of wind power. Croonenbroeck and Stadtmann (2015) show that this type of bias turns out to be very costly, from a turbine operator's point of view. GWPPT and the lasso, however, are still symmetric, mostly. For even longer forecasting horizons (Figures 8c and 8d show densities for one day and two days ahead forecasting errors), the AR asymmetry becomes worse, while GWPPT starts to return asymmetric forecasts as well. Also, the lasso model becomes asymmetric, but not as strongly as the other models. From that we conclude that using the lasso model instead of any of the other models may provide not only the most accurate forecasts, but also has the least severe impact of asymmetry, which is important for any turbine operator with respect to the financial impact of the forecast.

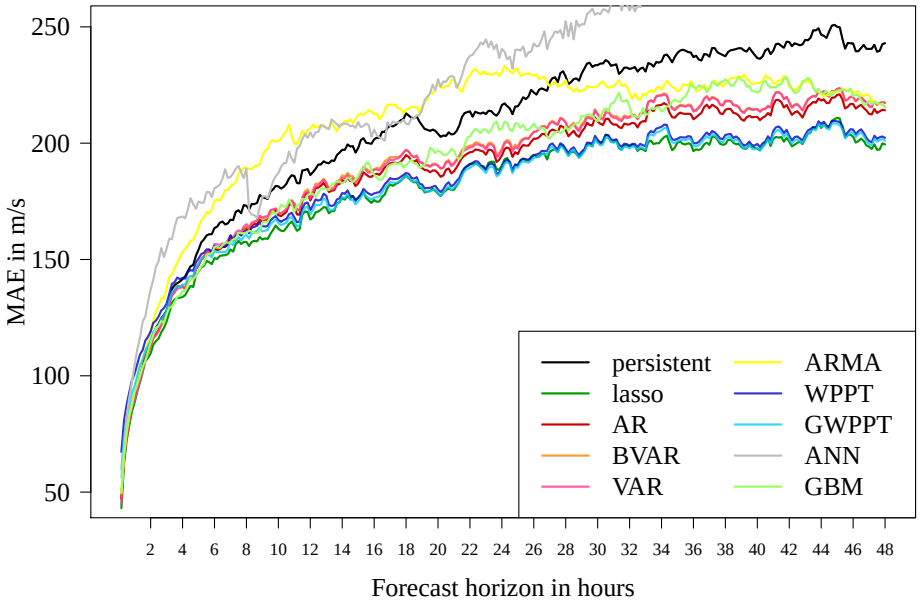

(a) MAE_k for the selected point forecasts.

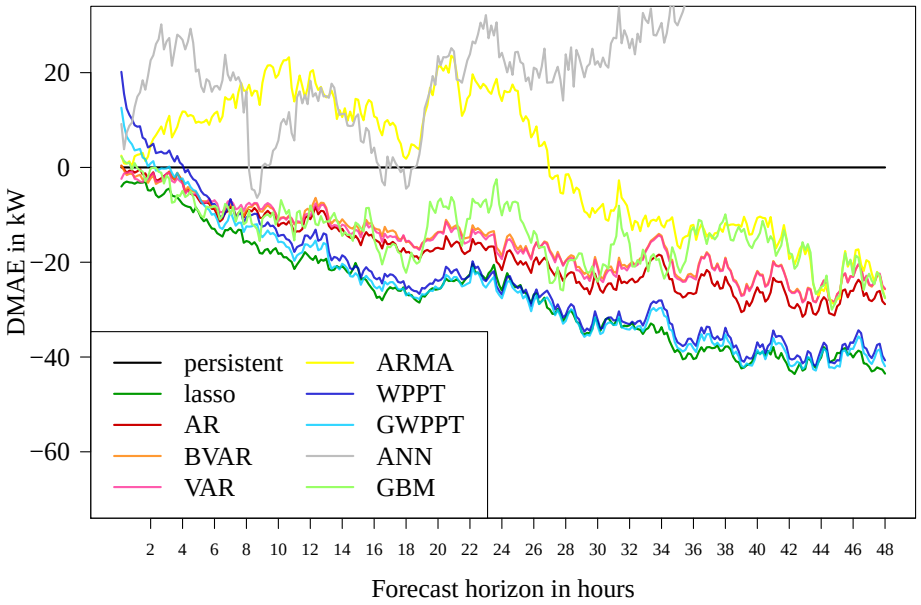

(b) DMAE_k (difference to persistence) for the selected point forecasts.

Figure 6: MAE_k and DMAE_k for all forecasting horizons k, time frame from November 2011 to November 2012.

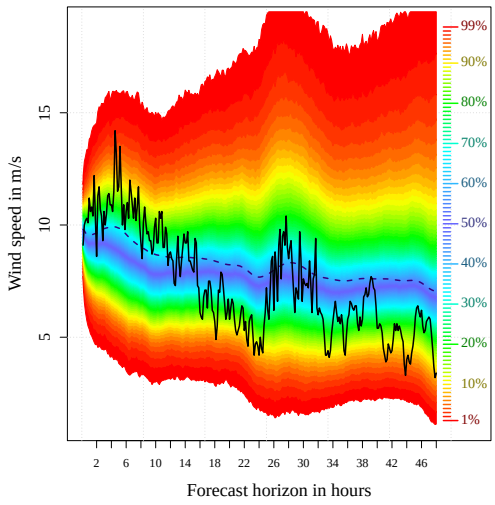

(a) Probabilistic wind speed forecast of Turbine A.

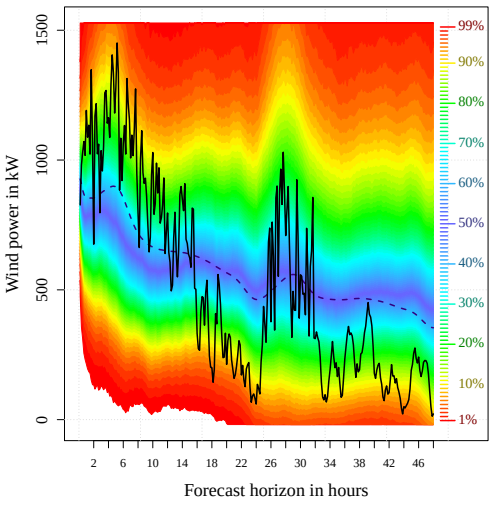

(b) Probabilistic wind power forecast of Turbine A.

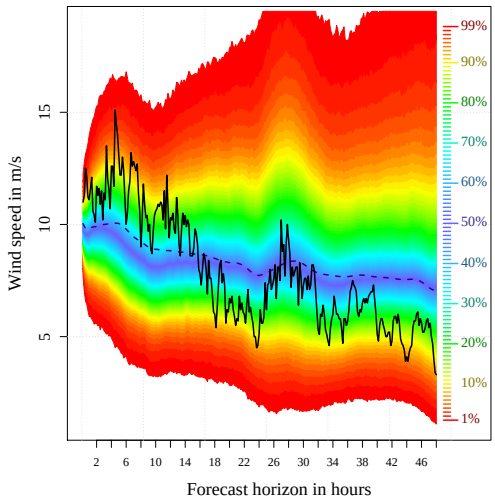

(c) Probabilistic wind speed forecast of Turbine B.

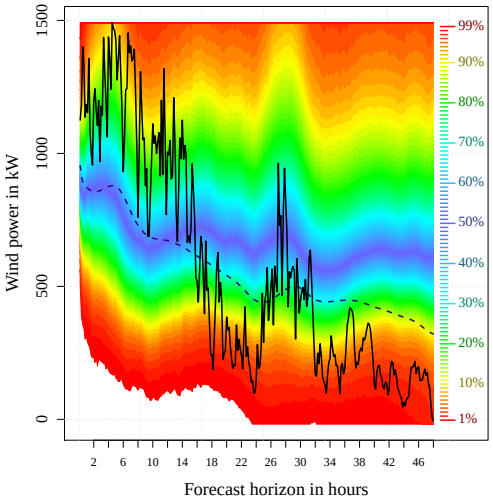

(d) Probabilistic wind power forecast of Turbine B.

Figure 7: Probabilistic wind speed and power forecast of Turbines A and B from 2012-02-25 07:20 to 2012-02-27 07:10. The black lines are the observed values, the dashed blue lines give the respective point estimates.

	1	6	24	48	72	144	288
persistent	47.12(0.61)	92.04(1.16)	141.58(1.59)	173.09(1.97)	186.30(2.19)	216.38(2.42)	242.99(2.59)
lasso	**43.10**(0.52)	**88.59**(1.08)	**133.87**(1.48)	**157.58**(1.79)	**167.20**(1.96)	190.42(2.19)	**199.49**(2.12)
AR	47.45(0.57)	90.93(1.03)	138.45(1.34)	163.73(1.58)	175.36(1.72)	194.72(1.88)	214.14(1.91)
BVAR	47.55(0.58)	89.99(1.03)	138.74(1.36)	164.10(1.58)	177.08(1.72)	197.69(1.89)	217.29(1.94)
VAR	44.72(0.50)	89.64(1.03)	138.19(1.35)	165.18(1.59)	176.02(1.69)	197.03(1.87)	217.40(1.93)
ARMA	49.61(0.53)	94.41(0.94)	153.39(1.44)	189.23(1.78)	206.41(1.87)	230.39(2.10)	215.40(1.84)
WPPT	67.30(0.78)	100.95(1.15)	142.09(1.52)	162.59(1.76)	171.26(1.90)	189.55(2.10)	202.28(2.17)
GWPPT	59.72(0.75)	95.64(1.15)	139.26(1.52)	160.66(1.76)	169.49(1.92)	**188.77**(2.12)	201.04(2.21)
ANN	56.29(0.59)	104.78(1.22)	168.50(1.76)	189.13(2.05)	204.61(2.06)	240.57(2.19)	296.65(2.03)
GBM	49.42(0.52)	92.96(0.93)	136.65(1.38)	161.55(1.74)	178.26(1.79)	206.27(2.10)	215.48(2.12)

Table 4: MAEs with estimated standard deviations. Best = bold, all within the 2-sigma range are underlined (not significantly worse than the best).

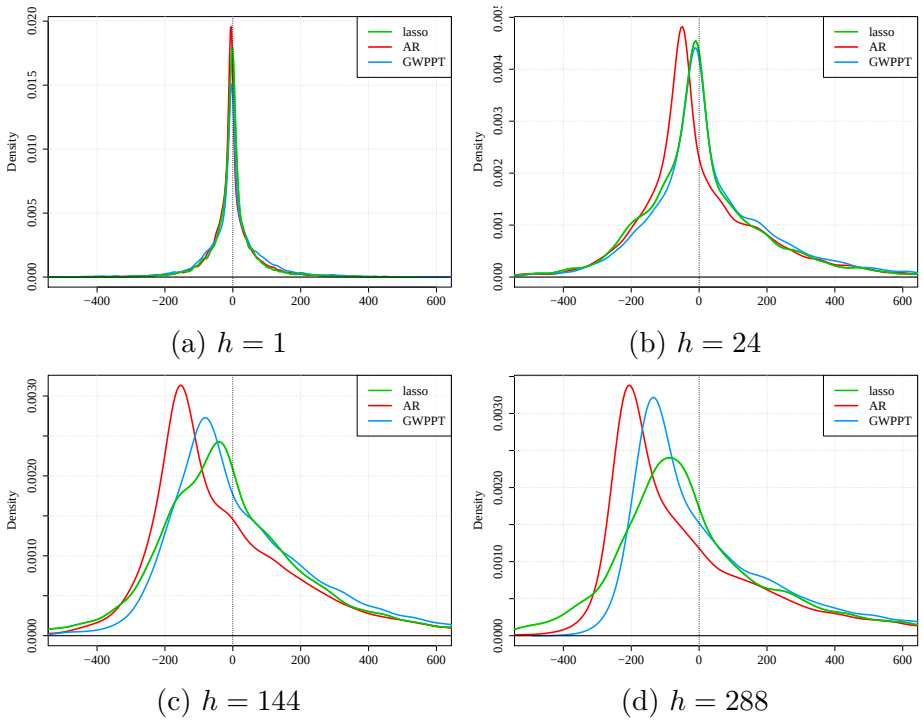

(a) $h = 1$ (b) $h = 24$

(c) $h = 144$ (d) $h = 288$

Figure 8: OOS forecasting errors density, lasso, AR and GWPPT, h step ahead forecasts, time frame from November 2011 to November 2012.

6. Conclusion

In this paper we present a new wind power forecasting approach that incorporates conditional heteroscedasticity, flexible periodicity and non-linearity modeling and provides the important wind speed forecasts in only one step. As an estimation technique, we present the re-weighted iterative lasso, which consumes little computing time, provides automatic regularization and sparsity and does not require a distributional assumption, unlike the usual maximum likelihood estimation.

The model for wind speed and wind power combines a multivariate time varying TVARMA process with a power-TGARCH model. The model returns accurate wind power forecasting results that are competitive, especially for the medium-term scenario. Furthermore, the model allows for probabilistic forecasting. Wind park operators may benefit from the minor asymmetry of

our model. While other models tend to over-estimate in increasing forecasting horizon settings, our model remains mostly stable, which helps keeping the asymmetry-induced financial loss of forecasts under control. For energy markets match-making, finally, our model provides not only superior accuracy for the point forecast necessary for both sellers and buyers, but also gives insight into the forecasts' distribution and thus, the forecasts' reliability.

Appendix A.

A B-spline basis function of degree H is constructed out of a B-spline basis function $\widetilde{B}$. $\widetilde{B}$ is defined by the degree H and a set of knots $\mathcal{K}$. The set of knot $\mathcal{K}$ contains $H+1$ knots $\{k_0, \ldots, k_{H+1}\}$ with $k_h < k_{h+1}$. This can be easily defined by the recurrence relation from (de Boor, 2001, p. 90):

$$\begin{aligned}
&\widetilde{B}(t; \{k_0, \ldots, k_{H+1}\}, H) \\
=&\frac{t - k_0}{k_H - k_0} \widetilde{B}(t; \{k_0, \ldots, k_H\}, H-1) \\
&+ \frac{t - k_1}{k_{H+1} - k_1} \widetilde{B}(t; \{k_1, \ldots, k_{H+1}\}, H-1)
\end{aligned} \tag{A.1}$$

with initialization

$$\widetilde{B}(t; \{k_l, k_{l+1}\}, 0) = \begin{cases} 1 & , t \in [k_l, k_{l+1}) \\ 0 & , \text{ otherwise.} \end{cases}$$

We consider the set of knots $\mathcal{K}(T, H)$ to be equidistant with center T. Thus, we find $k_0 = T - h\frac{D+1}{2}$, $k_{D+1} = T + h\frac{D+1}{2}$ and since we select an odd degree D, we get $k_{\frac{D+1}{2}} = T$, where h is the distance between the knots. Note that H and h define the knots $\mathcal{K}$ uniquely.
Finally, we consider a seasonality S to obtain a periodic basis function $\widetilde{B}(t; \mathcal{K}, H)$. To do so, it is suitable to choose h such that S is an integer multiple of h, which itself is at least $H+1$ to guarantee a partition of the unity. We define

$$\widetilde{B}_1^*(t; \mathcal{K}, H) = \sum_{k \in \mathbb{Z}} \widetilde{B}(t - kS; \mathcal{K}, H) \tag{A.2}$$

as the initial periodic basis function. In our setting, the data has two seasons, a diurnal and an annual one.[2] As our data frequency is at 10 minutes, we have six observations per hour. Thus, our diurnal seasons are $S_{\text{diurnal}} = 24 \times 6 = 144$ and the yearly seasons are $S_{\text{annual}} = 365.24 \times 24 \times 6 =$

[2]Wind speed as well as wind power can be assumed to be periodic for daily and yearly patterns. Empirically, this behavior can be shown by using periodograms, i.e. by analyzing the empirical spectral density.

52594.56.[3] By using the initial periodic basis function $\widetilde{B}_1^*$, we define the full periodic basis by $\widetilde{B}_j^*(t;\mathcal{K},H) = \widetilde{B}_{j-1}^*(t-h;\mathcal{K},H)$. In conclusion, the basis $\mathcal{B} = \{\widetilde{B}_1^*,\ldots,\widetilde{B}_{N_\mathcal{B}}^*\}$ has a total of $N_\mathcal{B} = S/h$ basis functions. In our setting, we choose $h_{\text{diurnal}} = 12$ and $h_{\text{annual}} = 4$.
The basis functions $\widetilde{B}_l^*$ are suitable to capture seasonal changes of parameters. However, due to the structure of the the basis functions, they model the absolute impact over time. In practice, it may be better to consider the changes over time instead of the absolute impact, especially if we use automatic shrinkage and selection algorithms for estimation, just as we do. We can easily model the changes in the parameters over time by cumulating the basis functions $\widetilde{B}_l^*$ within l. Hence, we define

$$\widetilde{B}_l^{*,\text{cum.}} = \widetilde{B}_{l-1}^{*,\text{cum.}} + \widetilde{B}_l^* \tag{A.3}$$

for $l > 1$ with $\widetilde{B}_1^{*,\text{cum.}} = \widetilde{B}_1^*$.
We use the cumulative basis functions for the conditional mean model (1), for both the diurnal and the annual basis functions. Also, we use the non-cumulative version for the conditional variance model (2), due to the parameter constraints in the variance model. We discuss this in greater detail in the estimation section.
As pointed out by Ziel et al. (2015), there might be interactions between the seasonal components. As the amount of sunshine is changing over the year, this might have impact on the wind speed. Thus, it is possible that daily cyclic effects are changing over the year. The simplest approach to model these interactions is to consider multiplications on the corresponding basis functions. We will use this multiplication for the conditional mean model, where we consider the cumulative basis functions. The multiplications are

$$\begin{aligned} B_{l_1 h_{\text{annual}}+l_2}^{\text{cum.}}(t) =& \widetilde{B}_{l_1}^{*,\text{cum.}}(t;\mathcal{K}(h_{\text{annual}},H),H)\times \\ & \widetilde{B}_{l_2}^{*,\text{cum.}}(t;\mathcal{K}(h_{\text{diurnal}},H),H) \end{aligned} \tag{A.4}$$

for $l_1 \in \{1,\ldots,h_{\text{diurnal}}\}$ and $l_2 \in \{1,\ldots,h_{\text{annual}}\}$ in equation (7) and for each periodic coefficient ξ in (1).

[3]Note that an average year lasts 365.242375 days, which is approximated by the leap year system every four years. A usual consensus is to approximate this by 365.24 days per year.

For the conditional variance equation, we do not consider the cumulative basis function, so here the multiplication is

$$\begin{aligned} B_{l_1 h_{\text{annual}}+l_2}(t) =& \widetilde{B}^*_{l_1}(t; \mathcal{K}(h_{\text{annual}}, H), H) \times \\ & \widetilde{B}^*_{l_2}(t; \mathcal{K}(h_{\text{diurnal}}, H), H) \end{aligned} \tag{A.5}$$

for $l_1 \in \{1, \ldots, h_{\text{diurnal}}\}$ and $l_2 \in \{1, \ldots, h_{\text{annual}}\}$ in equation (7) and for each periodic coefficient ξ of the conditional variance model (2).
However, by construction of the periodic basis, $\sum_{l=1}^{N_{\mathcal{B}}} \widetilde{B}^*_l(t)$ is constant. Thus, for the time varying coefficient ξ of the conditional mean model (1), we consider the set of basis functions

$$\begin{aligned} \mathcal{B}^{\text{cum.}}_{\xi} = \{ B^{\text{cum.}}_{l_1 h_{\text{annual}}+l_2} | & l_1 \in \{1, \ldots, h_{\text{diurnal}}\}, \\ & l_2 \in \{1, \ldots, h_{\text{annual}}\} \}, \end{aligned} \tag{A.6}$$

where the last element is constant. Note that $\mathcal{B}^{\text{cum.}}_{\xi}$ has $h_{\text{diurnal}} \times h_{\text{annual}}$ elements, so that in our setting, $12 \times 4 = 48$ parameters for each time varying coefficient.
For the conditional variance model (2), we define the used set of basis function for a periodic parameter ξ by

$$\begin{aligned} \mathcal{B}_{\xi} = \{1\} \cup \{ B_{l_1 h_{\text{annual}}+l_2} | & l_1 \in \{1, \ldots, h_{\text{diurnal}}\}, \\ & l_2 \in \{1, \ldots, h_{\text{annual}}\}, \\ & (l_1, l_2) \neq (1, 1) \}. \end{aligned} \tag{A.7}$$

Thus, we replace the first basis function B_1 by the constant 1, to model the constant impact directly.

References

Alessandrini, S., Sperati, S., Pinson, P., 2013. A comparison between the ecmwf and cosmo ensemble prediction systems applied to short-term wind power forecasting on real data. Applied Energy 107, 271–280.

Ambach, D., 2015. Short-term wind speed forecasting in germany. Journal of Applied Statistics, 1–19.

Ambach, D., Croonenbroeck, C., 2015. Space-time short-to medium-term wind speed forecasting. Statistical Methods & Applications, 1–16.

Ambach, D., Schmid, W., 2015. Periodic and long range dependent models for high frequency wind speed data. Energy 82, 277–293.

Amjady, N., Keynia, F., Zareipour, H., 2011. Short-term wind power forecasting using ridgelet neural network. Electric Power Systems Research 81 (12), 2099–2107.

Azad, H. B., Mekhilef, S., Ganapathy, V. G., 2014. Long-term wind speed forecasting and general pattern recognition using neural networks. Sustainable Energy, IEEE Transactions on 5 (2), 546–553.

Berkhout, V., Faulstich, S., Görg, P., Hahn, B., Linke, K., Neuschäfer, M., Pfaffel, S., Rafik, K., Rohrig, K., Rothkegel, R., Zieße, M., 2013. Wind energy report germany 2013. Fraunhofer-Institut für Windenergie und Energiesystemtechnik-IWES-Kassel.

Berner, J., Fossell, K. R., Ha, S.-Y., Hacker, J. P., Snyder, C., 2015. Increasing the skill of probabilistic forecasts: Understanding performance improvements from model-error representations. Monthly Weather Review 143 (4), 1295–1320.

Bhaskar, K., Singh, S., 2012. Awnn-assisted wind power forecasting using feed-forward neural network. Sustainable Energy, IEEE Transactions on 3 (2), 306–315.

Cadenas, E., Rivera, W., 2009. Short term wind speed forecasting in la venta, oaxaca, méxico, using artificial neural networks. Renewable Energy 34 (1), 274–278.

Cao, Q., Ewing, B. T., Thompson, M. A., 2012. Forecasting wind speed with recurrent neural networks. European Journal of Operational Research 221 (1), 148–154.

Carapellucci, R., Giordano, L., 2013. The effect of diurnal profile and seasonal wind regime on sizing grid-connected and off-grid wind power plants. Applied Energy 107, 364–376.

Chen, K., Chan, K.-S., 2011. Subset arma selection via the adaptive lasso. Statistics and its Interface 4 (2), 197–205.

Costa, A., Crespo, A., Navarro, J., Lizcano, G., Madsen, H., Feitosa, E., 2008. A review on the young history of the wind power short term prediction. Renewable and Sustainable Energy Reviews 12 (6), 1725–1744.

Croonenbroeck, C., Ambach, D., 2015. Censored spatial wind power prediction with random effects. Renewable and Sustainable Energy Reviews 51, 613–622.

Croonenbroeck, C., Dahl, C. M., 2014. Accurate medium-term wind power forecasting in a censored classification framework. Energy 73, 221 – 232.

Croonenbroeck, C., Stadtmann, G., 2015. Minimizing asymmetric loss in medium-term wind power forecasting. Renewable Energy 81, 197–208.

de Boor, C., 2001. A Practical Guide to Splines, revised Edition. Springer, New York.

De Giorgi, M. G., Ficarella, A., Tarantino, M., 2011. Error analysis of short term wind power prediction models. Applied Energy 88 (4), 1298–1311.

Erdem, E., Shi, J., 2011. Arma based approaches for forecasting the tuple of wind speed and direction. Applied Energy 88 (4), 1405–1414.

Evans, S. C., Zhang, Z., Iyengar, S., Chen, J., Hilton, J., Gregg, P., Eldridge, D., Jonkhof, M., McCulloch, C., Shokoohi-Yekta, M., 2014. Towards wind farm performance optimization through empirical models. In: Aerospace Conference, 2014 IEEE. IEEE, pp. 1–12.

Ewing, B. T., Kruse, J. B., Schroeder, J. L., 2006. Time series analysis of wind speed with time-varying turbulence. Environmetrics 17 (2), 119–127.

Friedman, J., Hastie, T., Höfling, H., Tibshirani, R., et al., 2007. Pathwise coordinate optimization. The Annals of Applied Statistics 1 (2), 302–332.

Giebel, G., Brownsword, R., Kariniotakis, G., Denhard, M., Draxl, C., 2011. The state-of-the-art in short-term prediction of wind power. Tech. rep., ANEMOS.plus, RisF8 DTU, Wind Energy Division.

Gneiting, T., Raftery, A. E., 2007. Strictly proper scoring rules, prediction, and estimation. Journal of the American Statistical Association 102, 359–378.

Haque, A. U., Nehrir, M. H., Mandal, P., 2014. A hybrid intelligent model for deterministic and quantile regression approach for probabilistic wind power forecasting. Power Systems, IEEE Transactions on 29 (4), 1663–1672.

Hennessey, J. P., 1977. Some aspects of wind power statistics. Journal of Applied Meteorology 16 (2), 119–128.

Hong, T., Pinson, P., Fan, S., Zareipour, H., Troccoli, A., Hyndman, R. J., 2016. Probabilistic energy forecasting: Global energy forecasting competition 2014 and beyond. International Journal of Forecasting, forthcoming.

Jeon, J., Taylor, J. W., 2012. Using conditional kernel density estimation for wind power density forecasting. Journal of the American Statistical Association 107 (497), 66–79.

Jung, J., Broadwater, R. P., 2014. Current status and future advances for wind speed and power forecasting. Renewable and Sustainable Energy Reviews 31, 762–777.

Kavasseri, R. G., Seetharaman, K., 2009. Day-ahead wind speed forecasting using f-arima models. Renewable Energy 34 (5), 1388–1393.

Landry, M., Erlinger, T. P., Patschke, D., Varrichio, C., 2016. Probabilistic gradient boosting machines for gefcom2014 wind forecasting. International Journal of Forecasting, forthcoming.

Lei, M., Shiyan, L., Chuanwen, J., Hongling, L., Zhang, Y., 2009. A review on the forecasting of wind speed and generated power. Renewable and Sustainable Energy Reviews 13, 915–920.

Li, G., Shi, J., 2010. On comparing three artificial neural networks for wind speed forecasting. Applied Energy 87 (7), 2313–2320.

Liu, D., Niu, D., Wang, H., Fan, L., 2014. Short-term wind speed forecasting using wavelet transform and support vector machines optimized by genetic algorithm. Renewable Energy 62, 592–597.

Newey, W. K., West, K. D., 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. Econometrica 55 (3), 703–708.

Nielsen, H. A., Pinson, P., Christiansen, L. E., Nielsen, T. S., Madsen, H., Badger, J., Giebel, G., Ravn, H. F., 2007. Improvement and automation of tools for short term wind power forecasting. Tech. rep., Scientific Proceedings of the European Wind Energy Conference & Exhibition, Milan, Italy.

Pinson, P., et al., 2013. Wind energy: Forecasting challenges for its operational management. Statistical Science 28 (4), 564–585.

Scholz, T., Lopes, V. V., Estanqueiro, A., 2014. A cyclic time-dependent markov process to model daily patterns in wind turbine power production. Energy 67, 557–568.

Shukur, O. B., Lee, M. H., 2015. Daily wind speed forecasting through hybrid kf-ann models based on arima. Renewable Energy 76, 637–647.

Silva, A. R., Pimenta, F. M., Assireu, A. T., Spyrides, M. H. C., 2016. Complementarity of brazils hydro and offshore wind power. Renewable and Sustainable Energy Reviews 56, 413–427.

Soman, S. S., Zareipour, H., Malik, O., Mandal, P., 2010. A review of wind power and wind speed forecasting methods with different time horizons. In: North American Power Symposium (NAPS), 2010. IEEE, pp. 1–8.

Tascikaraoglu, A., Uzunoglu, M., 2014. A review of combined approaches for prediction of short-term wind speed and power. Renewable and Sustainable Energy Reviews 34, 243–254.

Tibshirani, R., 1996. Regression shrinkage and selection via the lasso. Journal of the Royal Statistical Society. Series B (Methodological), 267–288.

Wagener, J., Dette, H., 2012. Bridge estimators and the adaptive lasso under heteroscedasticity. Mathematical Methods of Statistics 21 (2), 109–126.

Zeng, J., Qiao, W., 2012. Short-term wind power prediction using a wavelet support vector machine. Sustainable Energy, IEEE Transactions on 3 (2), 255–264.

Zhou, Z., Botterud, A., Wang, J., Bessa, R., Keko, H., Sumaili, J., Miranda, V., 2013. Application of probabilistic wind power forecasting in electricity markets. Wind Energy 16 (3), 321–338.

Zhu, X., Genton, M. G., Gu, Y., Xie, L., 2014. Space-time wind speed forecasting for improved power system dispatch. Test 23 (1), 1–25.

Ziel, F., 2015. Iteratively reweighted adaptive lasso for conditional heteroscedastic time series with applications to ar–arch type processes. Computational Statistics & Data Analysis, forthcoming.

Ziel, F., Liu, B., 2016. Lasso estimation for gefcom2014 probabilistic electric load forecasting. International Journal of Forecasting, forthcoming.

Ziel, F., Steinert, R., Husmann, S., 2015. Efficient modeling and forecasting of electricity spot prices. Energy Economics 47, 98–111.

Zugno, M., Jonsson, T., Pinson, P., 2012. Trading wind energy on the basis of probabilistic forecasts both of wind generation and of market quantities. Wind Energy 16 (6), 909–926.